

ROBUST MACHINE LEARNING APPLIED TO ASTRONOMICAL DATA SETS. II. QUANTIFYING PHOTOMETRIC REDSHIFTS FOR QUASARS USING INSTANCE-BASED LEARNING

NICHOLAS M. BALL,^{1,2} ROBERT J. BRUNNER,^{1,2} ADAM D. MYERS,^{1,2} NATALIE E. STRAND,³ STACEY L. ALBERTS,¹
 DAVID TCHENG,² AND XAVIER LLORÀ²

Received 2006 December 16; accepted 2007 March 16

ABSTRACT

We apply instance-based machine learning in the form of a k -nearest neighbor algorithm to the task of estimating photometric redshifts for 55,746 objects spectroscopically classified as quasars in the Fifth Data Release of the Sloan Digital Sky Survey. We compare the results obtained to those from an empirical color-redshift relation (CZR). In contrast to previously published results using CZRs, we find that the instance-based photometric redshifts are assigned with no regions of catastrophic failure. Remaining outliers are simply scattered about the ideal relation, in a manner similar to the pattern seen in the optical for normal galaxies at redshifts $z \lesssim 1$. The instance-based algorithm is trained on a representative sample of the data and pseudo-blind-tested on the remaining unseen data. The variance between the photometric and spectroscopic redshifts is $\sigma^2 = 0.123 \pm 0.002$ (compared to $\sigma^2 = 0.265 \pm 0.006$ for the CZR), and $54.9\% \pm 0.7\%$, $73.3\% \pm 0.6\%$, and $80.7\% \pm 0.3\%$ of the objects are within $\Delta z < 0.1$, 0.2 , and 0.3 , respectively. We also match our sample to the Second Data Release of the Galaxy Evolution Explorer legacy data, and the resulting 7642 objects show a further improvement, giving a variance of $\sigma^2 = 0.054 \pm 0.005$, with $70.8\% \pm 1.2\%$, $85.8\% \pm 1.0\%$, and $90.8\% \pm 0.7\%$ of objects within $\Delta z < 0.1$, 0.2 , and 0.3 . We show that the improvement is indeed due to the extra information provided by *GALEX*, by training on the same data set using purely SDSS photometry, which has a variance of $\sigma^2 = 0.090 \pm 0.007$. Each set of results represents a realistic standard for application to further data sets for which the spectra are representative.

Subject headings: catalogs — cosmology: miscellaneous — methods: data analysis — quasars: general

Online material: color figures

1. INTRODUCTION

Photometric redshifts, both from empirical training sets and template spectral energy distributions (SEDs), are important for the application of objects to the study of cosmology, as they enable the exploration of large regions of space that are otherwise inaccessible. This is achieved both in cosmological volume through a higher number density of objects and in parameter space through finer binning.

After the early work of Baum (1962), Koo (1985), and Loh & Spillar (1986), a variety of techniques were extensively developed (Gwyn & Hartwick 1996; Lanzetta et al. 1996; Mobasher et al. 1996; Sawicki et al. 1997; Connolly et al. 1998; Wang et al. 1998; Benítez 2000) on galaxies in the deep, but narrow, Hubble Deep Field North (HDF-N; Williams et al. 1996). These different methods were shown to be mutually consistent and relatively accurate in blind testing (Hogg et al. 1998).

More recently, wide-field surveys with multicolor photometry and fiber-based spectroscopy have generated large, uniform samples that enable photometric redshifts to be estimated for both galaxies and quasars. For galaxies in these surveys at redshifts of $z \lesssim 0.4$, (e.g., Brunner et al. 1997, 2000; Tagliaferri et al. 2002; Firth et al. 2003; Vanzella et al. 2004; Ball et al. 2004; Collister & Lahav 2004; Wadadekar 2005), a number of results have converged to an rms dispersion of $\sigma \sim 0.02$ (i.e., $\sigma^2 \sim 0.0004$) between spectroscopic and photometric redshifts, with no serious systematic effects. It should be emphasized, however, that galaxy photometry in these previous analyses has been very good, typically a few percent or better. Way & Srivastava (2006) show

similar results when combining the SDSS DR2 (Abazajian et al. 2004), *GALEX* GR1 (Martin et al. 2005) and the extended source catalog of the Two Micron All Sky Survey (Skrutskie et al. 2006). The results at moderate redshifts have also been successful, with luminous red galaxies (Eisenstein et al. 2001) in the SDSS trained with redshifts in the 2SLAQ survey (Cannon et al. 2006) having an rms of $\sigma = 0.049$ (Collister et al. 2007) for a sample at $0.4 < z < 0.7$ (see also Padmanabhan et al. 2005).

At high redshifts, the number of spectra available is smaller and, in addition to the HDF-N, there have been analyses of other deep fields, such as the HDF South (Williams et al. 2000) and the Hubble Ultra Deep Field (Beckwith et al. 2006). In the latter, Coe et al. (2006) show an accuracy of $\Delta z = 0.04(1+z)$ for $z \lesssim 6$.

In contrast to galaxies, which show small numbers of outliers, but no significant groups of outlying objects, all wide-field quasar photometric redshift results to date (Richards et al. 2001; Budavári et al. 2001; Weinstein et al. 2004, hereafter W04; Wu et al. 2004; Babbedge et al. 2004) suffer from regions of “catastrophic” failure, in which groups of objects are assigned a redshift very different from the true value. The first four use SDSS data, while the last uses the ELAIS N1 and N2 fields and the Chandra Deep Field North. W04 implement an empirical method based on color-redshift relations, which we use as our baseline. Catastrophic failures severely hamper cosmological investigations that use photometrically selected quasar samples (e.g., Myers et al. 2006, 2007a, 2007b), particularly by assigning objects at $z > 2$ to $z < 1$ and vice versa; thus, eliminating these regions is important. Reasons for the failures, depending on the details of the way a particular data set is chosen, include quasar reddening, degeneracy in the color-redshift relation, and superimposition of emission from another object, for example, an extended host galaxy.

Results using a more restricted parameter space (Wolf et al. 2003), defined by $17 < R < 24$ and $1.2 < z < 4.8$ in the 17 filter

¹ Department of Astronomy, University of Illinois, Urbana, IL 61801; nbball@astro.uiuc.edu.

² National Center for Supercomputing Applications, Urbana, IL 61801.

³ Department of Physics, University of Illinois, Urbana, IL 61801.

set of the COMBO-17 survey (e.g., Wolf et al. 2004), have met with more success. However, the sample size, 192 quasars, is small and limited in angular extent, and therefore is of limited cosmological applicability.

In this paper we use optical data from the Fifth Data Release of the SDSS and near- and far-UV data from the Second Data Release of the *Galaxy Evolution Explorer* (*GALEX*; Martin et al. 2005) to assign photometric redshifts to quasars. Our results improve upon previous wide-field techniques by eliminating regions of catastrophic failure, resulting in a distribution of quasar photometric redshifts comparable to those obtained for galaxies. We do not address the application of the photometric redshifts to any parameter space beyond that represented by the training and blind test sets.

2. DATA

We use data from the Fifth Data Release (DR5; SDSS collaboration, 2007, in preparation) of the Sloan Digital Sky Survey (SDSS; York et al. 2000) and the Second Data Release (GR2) of the *Galaxy Evolution Explorer* (Martin et al. 2005). We select primary nonrepeating observations of objects classified as quasars (`specClass = qso` or `hiz_qso`) in the `specObj` view of the SDSS DR5 Catalog Archive Server database. The `hiz_qso` objects are at redshifts of $z > 2.3$ and trigger the use of the $\text{Ly}\alpha$ finding code in the SDSS spectroscopic pipelines (J. A. Frieman et al. 2007, in preparation; D. J. Schlegel et al. 2007, in preparation). We also require that the spectroscopic flags `zWarning = 0` and `zStatus > 2`, and that all input magnitudes are not at clearly unphysical extreme values, being in the range 0–40. The resulting sample contains 55,746 quasars.

In addition to the SDSS sample, the SDSS objects are cross-matched to the primary photometric objects in the `photoObjAll` view of the *GALEX* GR2 database. We find 8174 matches within an RA + decl. tolerance of $4''$. Of these 532 have more than one match and the poorer match is rejected, leaving an SDSS+*GALEX* sample of 7642 unique matches. For the *GALEX* objects, we require `primary_flag = 1`, a detection in both near and far-UV bands, magnitudes again in the range 0–40, and the flags `fuv_artifact` and `nuv_artifact` to be 0. Throughout, the SDSS magnitudes are corrected for Galactic extinction using the dust maps of Schlegel et al. (1998) and the *GALEX* magnitudes using the $B - V$ (e_{-bv}) term inferred from these maps using the standard formula of Cardelli et al. (1989).

The resulting samples of 55,746 and 7642 objects form training sets used as input for the learning algorithms. The full set of object attributes for the SDSS sample consists of 16 training features. These are the colors $u - g$, $g - r$, $r - i$, and $i - z$, where the SDSS bands u , g , r , i , and z are given for each of the four magnitude types, PSF, fiber, Petrosian, and model (Stoughton et al. 2002). For SDSS+*GALEX*, we add the colors $fiuv - nuv$ and $nuv - u$, where u is given in each of the four SDSS magnitude types, resulting in 21 training features.

In addition to the SDSS and SDSS+*GALEX* data sets, we also analyze the SDSS+*GALEX* sample of objects, but using only SDSS features. This data set, referred to as “*GALEX*-SDSS-only,” enables us to quantify the level of improvement in SDSS+*GALEX* seen from the addition of the *GALEX*UV features, as opposed to possible improvement due to the sample only containing quasars that appear in both SDSS and *GALEX*.

3. ALGORITHMS

We implement instance-based learning on the SDSS, SDSS+*GALEX*, and *GALEX*-SDSS-only data sets. The results are compared to those on the same data for an empirical color-redshift

relation containing full probability density functions (N. E. Strand 2007, in preparation). We also study the utility of subsets of the full set of training features using genetic algorithms.

The machine learning is implemented in the Java environment Data-to-Knowledge (Welge et al. 2003). It is optimized through use of nationally peer-reviewed allocated time on the Xeon Linux cluster Tungsten at the National Center for Supercomputing Applications. This enables an extensive exploration of the parameter space describing the training features of the objects and the settings of the learning algorithms.

3.1. Instance-based Learning

Instance-based learning (IB; e.g., Aha et al. 1991; Witten & Frank 2000; Hastie et al. 2001), is a powerful class of empirical machine learning methods that to date has not been extensively utilized on large astronomical data sets due to its computational intensity. Two examples where the method has been used are Budavári et al. (2001) and Csabai et al. (2003), who both applied it to the SDSS Early Data Release (EDR; Stoughton et al. 2002). However, they only utilize single nearest neighbors, and in addition the DR5 data set analyzed here is approximately 15 times the size of the EDR. Here, through the use of Tungsten (§ 3), we are able to realize the full potential of the algorithm, via the use of the k -nearest neighbor method (e.g., Cover & Hart 1967).

In its simplest form, the “training” of the algorithm is trivial and involves simply memorizing the positions of each of the examples in the training set. For each object in the testing set, the nearest training example is then found, and the predicted value, either a classification or a continuous value, is taken to be that of the training example. Thus, the computational expense is incurred at the time of classification, as a large number of distance calculations must be performed. However, the method is powerful because it uses all of the information available in the training set, rather than a model of the training set as is typically used by most other learning algorithms.

There are a number of simple refinements to this method, which in practice result in large improvements in performance: (1) Instead of the nearest neighbor to the testing example, the k nearest neighbors can be found, and the distances can be weighted using a predictive integration function to produce a weighted output. This function, d , takes the form

$$d = \sum_i^k \frac{1}{x_i^p},$$

where the x_i are the Euclidean distances to the neighbors and the exponent p can take on any positive value, typically, but not necessarily, an integer. (2) The input features can be *standardized* such that the mean and variance of each are 0 and 1, respectively. This prevents the training from being dominated by features with larger numerical values or spreads. Alternatively, one could also normalize the range of features to be 0–1. (3) Objects in the training set can be allocated to collective regions of parameter space, which can considerably reduce the required number of distance calculations.

Of the methods described, we implement (1) and (2), but not (3), as we wish to use the full information available in the training data. We optimize the values of k and p and standardize all training features. Further refinements can also be made for objects that have noncontinuous values, such as a classification or missing data. However, in this paper all values are considered; i.e., the training features and the spectroscopic and photometric redshifts are continuous.

3.2. Color-Redshift Relation

We have implemented the color-redshift relation (CZR) method of W04 on the same data as the IB. This enables a direct comparison of the performance of the two methods. The CZR establishes an empirical relation between the spectroscopic redshifts and the colors of the training set. The maximum likelihood redshift probability density function (PDF) is then found for each object in the test set.

3.3. Genetic Algorithms

The methods above select and optimize a learning algorithm for a given set of training features. However, it is possible that different subsets of the features available will produce better results. In particular, the results for instance-based learning can be made worse by noise in the training set or by irrelevant training features. To explore this possibility, we implement a binary genetic algorithm on the training feature sets.

A genetic algorithm (GA; e.g., Holland 1975; Goldberg 1989; Haupt & Haupt 1998) mimics evolution, in the sense that the most successful individuals are those that are best adapted for the task at hand. We implement the binary genetic algorithm, in which each individual is a string of 0's and 1's, which represents whether or not to use a particular input feature (in our case, the 16 colors). An initial population of random individuals is created, and the IB is run using the features selected. The result, in this case the variance between photometric and spectroscopic redshift, is the *fitness* of that individual. The individuals and their fitnesses are then combined to produce new individuals, and those with higher fitnesses are favored. In principle, a good approximation to the best set of features to use as the training set should be selected with this approach.

The combination involves identifying the best individuals to breed via tournament selection, in which a specified number of individuals from the population are selected and the best is put in the mating pool to be combined with other individuals. Two individuals are combined using one-point crossover, in which a segment of one is swapped with that of the other. To more fully explore the parameter space and prevent the algorithm from converging too rapidly on a local minimum, a probability of mutation is introduced on the newly created individuals before they are processed. This is simply the probability that a 0 becomes a 1, or vice versa.

An approximate number of individuals to use is given by

$$n_{\text{in}} \sim 2n_f \log n_f,$$

where n_f is the number of features. For the SDSS and *GALEX*-SDSS-only, $n_f = 16$ and for SDSS+*GALEX*, $n_f = 21$. Hence, $n_{\text{in}} \sim 39$ and 56, respectively, for these two values of n_f . The algorithm converges, i.e., finds the best individual and hence the best training set, in

$$n_{\text{it}} \sim \alpha n_f \log n_f$$

iterations, where α is a problem-dependent constant. Generally, $\alpha > 3$, giving an expected value for our data of $n_{\text{it}} \sim 58$ for $n_f = 16$ and $n_{\text{it}} \sim 83$ for $n_f = 21$. We employ this number of iterations with larger numbers of individuals⁴ to be sure that the algorithm has converged. Further information on genetic algorithm design can be found, e.g., in Goldberg (2002).

⁴ For SDSS+*GALEX*, 300 and for the other two data sets, 200. These numbers were selected for other tests not reported here and simply strengthen the null result.

Our GA is implemented on the IB for each of the SDSS, SDSS+*GALEX*, and *GALEX*-SDSS-only data sets. The settings of these algorithms are fixed for the duration of the GA iteration. It is possible in principle to combine the optimization of the learning algorithm and the feature set; however, we defer this analysis to a later paper.

3.4. Training and Quality of Redshifts

The IB and CZR are supervised learning algorithms; they are given a training set of objects and attempt to minimize a cost function that describes the quality of the predictions on a separate testing set. For IB, the cost function is given by the variance between the photometric and spectroscopic redshifts for objects with spectra,

$$\langle (\Delta z)^2 \rangle - \langle \Delta z \rangle^2,$$

where $\Delta z = |z_{\text{spec}} - z_{\text{phot}}|$, z_{spec} is the spectroscopic redshift value, and z_{phot} is the photometric redshift prediction made by the learning algorithm. The second term in the variance equation is small.

The value of the variance is dominated by the outliers. However, in our case, this is a desirable property, because it is these objects that we wish to pull in the most toward the correct values. The dominance of the outliers renders the variance susceptible to variations in this population. We therefore quote errors on all of our blind test variances, derived from splitting the population using multiple random seeds (see below). For the CZR, the cost function is the likelihood of the PDF.

Instance-based learning, like any supervised machine learning algorithm, is susceptible to incompleteness and noise in the training set. At present, the SDSS DR5 is by far the largest and most homogeneous quasar data set available, and it has a high completeness (e.g., Vanden Berk et al. 2005). Other available data sets are either not as deep, smaller (e.g., Croom et al. 2004), or deeper but orders of magnitude smaller (e.g., Wolf et al. 2004). One could prune noisy exemplars; however, it is difficult to meaningfully define what makes a noisy or sparsely populated region of parameter space, and pruning particular regions could introduce new and poorly defined biases. The use of multiple nearest neighbors smooths the noise, and the blind test results address both incompleteness and noise by presenting realistic results on unseen data.

The distance measure parameters of a number of nearest neighbors and the distance weighting assume that the input training features are uncorrelated; however, given that we repeat the same four colors in four magnitude types and that a set of features is always derived from a particular object, the input features will always be correlated, both in magnitude type (e.g., PSF $u - g$ is correlated to fiber $u - g$, and so on) and in color (e.g., PSF $u - g$ is correlated to PSF $g - r$, and so on.) Correlated input features are therefore unavoidable; we feel, however, that our algorithmic approach is acceptable, because we select the parameters to produce the optimal blind test result.

Different splits of the training set are investigated at various points in the learning process, giving four adjustable ratios: (1) The quantity r_{train} is the ratio between the data used as the training set and for testing the algorithm's performance according to the cost function to adjust the final model settings (for IB there is no adjustment, so the ratio just affects the performance through the information available). (2) The quantity r_{blind} is the ratio of the whole set of data used in training and testing to that unseen by the algorithm until it is applied, as it would be to new data

TABLE 1
SUMMARY OF PHOTOMETRIC REDSHIFT SAMPLES DESCRIBED IN THIS PAPER

Data Set	Method	Variance	Variance/(1 + z)	Mean $\Delta z/(1 + z)$	$\Delta z < 0.1$ (%)	$\Delta z < 0.2$ (%)	$\Delta z < 0.3$ (%)
SDSS	IB	0.123 ± 0.002	0.034 ± 0.001	0.095 ± 0.001	54.9 ± 0.7	73.3 ± 0.6	80.7 ± 0.3
SDSS+GALEX	IB	0.054 ± 0.005	0.014 ± 0.002	0.060 ± 0.003	70.8 ± 1.2	85.8 ± 1.0	90.8 ± 0.7
GALEX-SDSS-only	IB	0.090 ± 0.007	0.022 ± 0.001	0.081 ± 0.003	62.0 ± 1.4	78.9 ± 1.0	85.2 ± 1.2
SDSS	CZR	0.265 ± 0.006	0.079 ± 0.003	0.115 ± 0.002	63.9 ± 0.3	80.2 ± 0.4	85.7 ± 0.3
SDSS+GALEX	CZR	0.136 ± 0.015	0.031 ± 0.006	0.071 ± 0.005	74.9 ± 1.4	86.9 ± 0.6	91.0 ± 0.8
GALEX-SDSS-only	CZR	0.158 ± 0.013	0.041 ± 0.004	0.081 ± 0.004	74.1 ± 0.8	86.2 ± 0.7	89.7 ± 0.6

from another survey; this is the pseudoblind test. (3) The quantity r_{bag} is the ratio of the data used in each bagged model to the rest of the training data, where the training data is r_{train} of the whole data set. (4) The quantity r_{cross} is similar, but for cross-validation. The latter is distinguished from bagging because it takes different random subsamples of the whole r_{train} training and $1 - r_{\text{train}}$ testing set, whereas bagging subsamples r_{train} .

The value for which we quote results for all of these ratios is 80 : 20. For application to new data not used here, the value of r_{train} would be 100%, to maximize the information available. This is the standard σ^2 reported in the literature for CZR techniques, but its value would be meaningless for instance-based approaches.

For IB, the variances obtained are quoted from the pseudoblind test, as this represents the most realistic standard of performance available from within the SDSS and GALEX data sets to be expected on new data. The value of r_{blind} is always such that the training data are representative of the full data set.

We quote the mean and standard deviation of the best variance from 10 training runs with differing random seeds for r_{blind} . Each run produces a grid of models with the range $1 \leq k \leq 50$ and $1 \leq p \leq 10$, where k is the number of nearest neighbors and p is the exponent in the distance-weighting function (§ 3.1). Integral values of k and p were used, although this is not a requirement. We use positive values of p , as negative values would result in objects other than the nearest neighbor being given the highest weighting, which would be unphysical, as increasingly large values of k would be given an ever higher weight. We investigated bagging and cross-validation using values of r_{bag} and $r_{\text{cross-val}}$ of 80 : 20 and 50 : 50, but these were not found to be necessary for IB. Other measures, such as $\Delta z/(1 + z)$ and the percentage of objects within $\Delta z < 0.1, 0.2$, and 0.3 , are also given for comparison to other work. We do not quote any results in which there is any overlap between the training and testing data.

The comparative CZR results were obtained by using a 10-fold bootstrapped pseudoblind test, again in the ratio $r_{\text{blind}} = 80 : 20$.

4. RESULTS

We now describe results for the full SDSS DR5, SDSS DR5+GALEX GR2, and GALEX-SDSS-only data sets, all of which are summarized in Table 1.

4.1. SDSS DR5

We found that the ideal parameters are 22 ± 5 nearest neighbors (NN) and a distance weighting (DW) of 3.7 ± 0.5 . In the pseudoblind test on the unseen 20% of the data, the best variance between the photometric and spectroscopic redshifts is 0.123 ± 0.002 . A comparison between the photometric and spectroscopic redshifts is shown in Figure 1, and the effect of varying the NN and the DW for the pseudoblind test is shown in Figure 2. We find that $54.9\% \pm 0.7\%$, $73.3\% \pm 0.6\%$, and $80.7\% \pm 0.3\%$

of the objects are within $\Delta z < 0.1, 0.2$, and 0.3 , respectively. The variance weighted by redshift is $\sigma_z^2 = 0.034 \pm 0.001$, and the mean $\Delta z/(1 + z) = 0.095 \pm 0.001$.

Because the values of NN and DW used here are discrete (in principle, they can be continuous, but that was not attempted), the results presented in Figure 1 were obtained with the values of NN, DW, and the blind test set random seed that gave the best variance in its grid that was closest to the mean. Here these values are NN = 22, DW = 4, and a random seed of 8 (for the seeds we used the integers 0–9). The variance is 0.1240, which is consistent with the mean variance quoted.

Our key result, shown in Figure 1, is the absence of regions of catastrophic failure—there is no upturn in a histogram of Δz values at large Δz , just a smooth decline such that few objects are outliers. This is in contrast to previous results for quasar photometric redshifts, which, while showing a comparable spread of objects with low Δz , show outlying regions of objects




FIG. 1.—Contour plot of quasar photometric redshifts assigned by the instance-based learner vs. spectroscopic redshifts for the SDSS DR5 pseudoblind testing sample of 11,149 of 55,746 quasars described in the text. For contouring, the objects are placed in bins of 0.05 in redshift, although the values on both axes are continuous. The variance between the two measures over the whole redshift range is 0.123 ± 0.002 . Compared to Fig. 3, there are no regions of “catastrophic” failure, in which objects are assigned a redshift very different from the true value, just a smoothly declining spread of outliers. There are no objects outside the range of redshifts plotted. [See the electronic edition of the Journal for a color version of this figure.]




FIG. 2.—Effect of varying the number of nearest neighbors (NN) and the distance weighting (DW) of the instance-based learner for the pseudoblind test on the SDSS DR5 data set, showing the mean from 10 different training to pseudoblind test splits of the data with a varying random seed. The model that gives the lowest variance is marked with 1σ error bars. [See the electronic edition of the Journal for a color version of this figure.]

with high Δz . The scattering of outliers obtained by the IB is similar in form to that seen in other studies for normal galaxies at redshifts of $z \lesssim 1$ (see, for example, Fig. 3 of Ball et al. [2004] for SDSS Main Sample galaxies, which have a mean redshift of $z \sim 0.1$), although there is still structure seen in Figure 1, especially at $z_{\text{spec}} \lesssim 1$ and ~ 2.2 .

We have also implemented the methods of W04 on the SDSS DR3, without removing the reddened quasars (N. E. Strand 2007, in preparation). Here we apply that method to the SDSS DR5 data set as a direct comparison between the empirical CZR and the IB. We find that the CZR has slightly narrower dispersion than the IB, with Δz percentages of $63.9\% \pm 0.3\%$, $80.2\% \pm 0.4\%$, and $85.7\% \pm 0.3\%$ within $\Delta z < 0.1$, 0.2 , and 0.3 . However, as shown in Figure 3, it still shows regions of catastrophic failure. The variance is therefore significantly higher, at $\sigma^2 = 0.265 \pm 0.006$. We again plot the run from the 10 with the closest variance to the mean. In this case, this was the final run of the 10, with $\sigma^2 = 0.2653$.

Previous results using empirical CZRs show a similar pattern. For example, Figure 4 of W04 shows regions of quasars at $0 \lesssim z_{\text{phot}} \lesssim 1$ and $1.5 \lesssim z_{\text{phot}} \lesssim 4.5$ over the spectroscopic redshift range $0 \lesssim z_{\text{spec}} \lesssim 4$. Similar results are seen in Budavári et al. (2001), Richards et al. (2001), and Wu et al. (2004).

4.2. SDSS DR5+GALEX GR2

Adding the GALEX data significantly improves the results, as shown in Figures 4 and 5. Here we obtain a variance of 0.054 ± 0.005 for the pseudoblind test, $70.8\% \pm 1.2\%$, $85.8\% \pm 1.0\%$, and $90.8\% \pm 0.7\%$ of objects within $\Delta z < 0.1$, 0.2 , and 0.3 , $\sigma_z^2 = 0.014 \pm 0.002$, and the mean $\Delta z/(1+z) = 0.060 \pm 0.003$.

The number of nearest neighbors and distance weighting are 17 ± 5 and 4.4 ± 0.8 , respectively. A higher distance weighting is expected due to the greater dimensionality of the training feature space (21 colors instead of 16) compared to the SDSS




FIG. 3.—As in Fig. 1, but showing the results for the CZR photometric redshifts. The regions of catastrophic failure are seen, and the overall variance is $\sigma^2 = 0.265 \pm 0.006$. The values of z_{phot} resulting from this method are in bins of width 0.05. Here a uniformly distributed random offset up to ± 0.025 has been added to the values of z_{phot} for clarity. [See the electronic edition of the Journal for a color version of this figure.]

data set. The exact values of NN and DW that are plotted in Figure 4 are chosen in the same manner as for the SDSS and are $NN = 12$, $DW = 5$ and a random seed of 3. The variance is 0.0521 .

To show that the improvement is not simply due to the smaller set of objects that appear in both surveys (for example,




FIG. 4.—As in Fig. 1, but showing the results for 1528 of 7642 quasars present in the SDSS DR5 cross-matched to the GALEX GR2. The variance is improved to $\sigma^2 = 0.054 \pm 0.005$. [See the electronic edition of the Journal for a color version of this figure.]




FIG. 5.—As in Fig. 2, but for the SDSS+GALEX data set shown in Fig. 4. [See the electronic edition of the Journal for a color version of this figure.]

these objects may be brighter quasars in the SDSS with better photometry), we also applied the SDSS training procedure to the cross-matched sample. This gives better results than the SDSS sample, but they are still significantly worse than SDSS+GALEX. The variance is $\sigma^2 = 0.090 \pm 0.007$, and the other results are as seen in Table 1.

The SDSS results extend deeper than those matched with GALEX, to $z \lesssim 6$ rather than $z \lesssim 3.5$. The lack of quasars in the “redshift desert” at $z \gtrsim 2.2$ is seen in Figure 4, caused by the Lyman break in the spectrum at a rest-frame wavelength of 912 Å being shifted out of the UV.

The CZR results for SDSS+GALEX also improve over those from the full SDSS data set; $74.9\% \pm 1.4\%$, $86.9\% \pm 0.6\%$, and $91.0\% \pm 0.8\%$ of the objects are within $\Delta z < 0.1$, 0.2 , and 0.3 , respectively. This is still slightly better than IB for $\Delta z < 0.1$ and $\Delta z < 0.2$, but is the same for $\Delta z < 0.3$.

4.3. Genetic Algorithms

The application of the genetic algorithms to the SDSS, SDSS+GALEX, and GALEX-SDSS-only data sets converged on the use of approximately half of the training parameters, but the variance was not significantly different from that from using the full set of training features. The full sets were therefore used throughout. The result indicates that there is some redundancy in the training features, which is expected, given that they are measuring the four colors four different times, just through different apertures.

5. DISCUSSION

Although the results here represent an important step in the sense that there are no regions of catastrophic failure, further improvement is still possible. In particular: (1) The input object parameter distributions may be generalized into the form of a PDF for each object, which can be propagated through the learning process, to make more explicit those objects for which the redshift is less certain, to take into account the error on each parameter, and to output a PDF for each object instead of a scalar value. (2) The

noncatastrophics of the instance-based and the lower dispersion at low values of Δz of the CZR can be combined into a new learning algorithm. The IB is in fact able to obtain results similar to the CZR (i.e., an approximately 5% narrower dispersion and regions of catastrophic failure instead of a spread of objects) by using the single nearest neighbor instead of k nearest neighbors. (3) The addition of other multiwavelength training data, such as infrared data from UKIDSS (Lawrence et al. 2006) and *Spitzer* (Werner et al. 2004), can be included in the training process.

We also obtained quasar photometric redshifts using decision trees, as used in Ball et al. (2006) for star-galaxy separation. The variances obtained were generally comparable to, but slightly worse than, those for instance-based, and are, therefore, not reported here.

6. CONCLUSIONS

We apply instance-based machine learning to 55,746 objects spectroscopically classified as quasars in the Fifth Data Release of the Sloan Digital Sky Survey (SDSS) and to 7642 objects cross-matched from this sample to the Second Data Release of the *Galaxy Evolution Explorer* legacy data (SDSS+GALEX). The algorithm is able to assign photometric redshifts to quasars without regions of catastrophic failure, unlike previously published results. This will enable samples of quasars to be constructed for cosmological studies with minimal contamination from objects at severely incorrect redshifts. We obtain, for the same data, empirical color-redshift relations with full probability distributions and find that these are similar to previous results in the literature.

For SDSS, we find a photometric to spectroscopic variance of 0.123 ± 0.002 for a sample of the data not used in the training. For SDSS+GALEX, this improves to 0.054 ± 0.005 . Using purely SDSS on the latter data set (GALEX-SDSS-only), the variance is 0.090 ± 0.007 . Hence the improvement results from the extra UV information provided by GALEX and not the reduced sample size, better photometry, or lower redshifts. The percentages of objects within $\Delta z < 0.1$ are $54.9\% \pm 0.7\%$, $70.8\% \pm 1.2\%$, and $62.0\% \pm 1.4\%$ for SDSS, SDSS+GALEX, and GALEX-SDSS-only, respectively. Each set of results represents a realistic standard for application to further data sets of which the spectra are representative.

We thank the referee for a prompt and useful report, which improved the paper, and Kumara Sastry of the Illinois Genetic Algorithms Laboratory for a clarification of our use of specific genetic algorithms. The authors acknowledge support from NASA through grants NN6066H156 and 05-GALEX05-0036, from Microsoft Research, and from the University of Illinois. The authors made extensive use of the storage and computing facilities at the National Center for Supercomputing Applications and thank the technical staff for their assistance in enabling this work.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the US Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web site is <http://www.sdss.org>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical

Institute Potsdam, the University of Basel, Cambridge University, Case Western Reserve University, the University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPA), the Max-Planck-Institute for Astrophysics (MPIA), New Mexico State University, Ohio State University, the University of Pittsburgh, the University of Portsmouth, Princeton University,

the United States Naval Observatory, and the University of Washington.

Based on observations made with the NASA Galaxy Evolution Explorer. *GALEX* is operated for NASA by the California Institute of Technology under NASA contract NAS5-98034.

Data To Knowledge (D2K) software, D2K modules, and/or D2K itineraries used by us were developed at the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign.

This research has made use of NASA's Astrophysics Data System.

REFERENCES

- Abazajian, K., et al. 2004, *AJ*, 128, 502
- Aha, D. W., Kibler, D., & Albert, M. K. 1991, *Machine Learning*, 6, 37
- Babbedge, T. S. R., et al. 2004, *MNRAS*, 353, 654
- Ball, N. M., Brunner, R. J., Myers, A. D., & Tchong, D. 2006, *ApJ*, 650, 497
- Ball, N. M., Loveday, J., Fukugita, M., Nakamura, O., Okamura, S., Brinkmann, J., & Brunner, R. J. 2004, *MNRAS*, 348, 1038
- Baum, W. A. 1962, in *IAU Symp. 15, Problems of Extra-Galactic Research*, ed. G. C. McVittie (New York: Macmillan), 390
- Beckwith, S. V. W., et al. 2006, *AJ*, 132, 1729
- Benítez, N. 2000, *ApJ*, 536, 571
- Brunner, R. J., Connolly, A. J., Szalay, A. S., & Bershad, M. A. 1997, *ApJ*, 482, L21
- Brunner, R. J., Szalay, A. S., & Connolly, A. J. 2000, *ApJ*, 541, 527
- Budavári, T., et al. 2001, *AJ*, 122, 1163
- Cannon, R., et al. 2006, *MNRAS*, 372, 425
- Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, 345, 245
- Coe, D., Benítez, N., Sánchez, S. F., Jee, M., Bouwens, R., & Ford, H. 2006, *AJ*, 132, 926
- Collister, A., et al. 2007, *MNRAS*, 375, 68
- Collister, A. A., & Lahav, O. 2004, *PASP*, 116, 345
- Connolly, A. J., Szalay, A. S., & Brunner, R. J. 1998, *ApJ*, 499, L125
- Cover, T. M., & Hart, P. E. 1967, *IEEE Trans. Inf. Theory*, 13, 21
- Croom, S. M., Smith, R. J., Boyle, B. J., Shanks, T., Miller, L., Outram, P. J., & Loaring, N. S. 2004, *MNRAS*, 349, 1397
- Csabai, I., et al. 2003, *AJ*, 125, 580
- Eisenstein, D. J., et al. 2001, *AJ*, 122, 2267
- Firth, A. E., Lahav, O., & Somerville, R. S. 2003, *MNRAS*, 339, 1195
- Goldberg, D. E. 1989, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Reading: Addison Wesley)
- . 2002, *Design of Innovation: Lessons from and for Competent Genetic Algorithms* (Boston: Kluwer)
- Gwyn, S. D. J., & Hartwick, F. D. A. 1996, *ApJ*, 468, L77
- Hastie, T., Tibshirani, R., & Friedman, J. 2001, *The Elements of Statistical Learning* (Springer)
- Haupt, R. L., & Haupt, S. E. 1998, *Practical Genetic Algorithms* (New York: Wiley)
- Hogg, D. W., et al. 1998, *AJ*, 115, 1418
- Holland, J. H. 1975, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (Ann Arbor: Univ. Michigan Press)
- Koo, D. C. 1985, *AJ*, 90, 418
- Lanzetta, K. M., Yahil, A., & Fernandez-Soto, A. 1996, *Nature*, 381, 759
- Lawrence, A., et al. 2006, *MNRAS*, submitted (astro-ph/0604426)
- Loh, E. D., & Spillar, E. J. 1986, *ApJ*, 303, 154
- Martin, D. C., et al. 2005, *ApJ*, 619, L1
- Mobasher, B., Rowan-Robinson, M., Georgakakis, A., & Eaton, N. 1996, *MNRAS*, 282, L7
- Myers, A. D., Brunner, R. J., Nichol, R. C., Richards, G. T., Schneider, D. P., & Bahcall, N. A. 2007a, *ApJ*, 658, 85
- Myers, A. D., Brunner, R. J., Richards, G. T., Nichol, R. C., Schneider, D. P., & Bahcall, N. A. 2007b, *ApJ*, 658, 99
- Myers, A. D., et al. 2006, *ApJ*, 638, 622
- Padmanabhan, N., et al. 2005, *MNRAS*, 359, 237
- Richards, G. T., et al. 2001, *AJ*, 122, 1151
- Sawicki, M. J., Lin, H., & Yee, H. K. C. 1997, *AJ*, 113, 1
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Skrutskie, M. F., et al. 2006, *AJ*, 131, 1163
- Stoughton, C., et al. 2002, *AJ*, 123, 485
- Tagliaferri, R., Longo, G., Andreon, S., Capozziello, S., Donalek, C., & Giordano, G. 2002, preprint (astro-ph/0203445)
- Vanden Berk, D. E., et al. 2005, *AJ*, 129, 2047
- Vanzella, E., et al. 2004, *A&A*, 423, 761
- Wadadekar, Y. 2005, *PASP*, 117, 79
- Wang, Y., Bahcall, N., & Turner, E. L. 1998, *AJ*, 116, 2081
- Way, M. J., & Srivastava, A. N. 2006, *ApJ*, 647, 102
- Weinstein, M. A., et al. 2004, *ApJS*, 155, 243 (W04)
- Welge, M., et al. 2003, *Data to Knowledge* (Urbana: NCSA), <http://alg.ncsa.uiuc.edu/do/tools/d2k>.
- Werner, M. W., et al. 2004, *ApJS*, 154, 1
- Williams, R. E., et al. 1996, *AJ*, 112, 1335
- . 2000, *AJ*, 120, 2735
- Witten, I. H., & Frank, E. 2000, *Data Mining* (San Francisco: Morgan Kaufmann)
- Wolf, C., Wisotzki, L., Borch, A., Dye, S., Kleinheinrich, M., & Meisenheimer, K. 2003, *A&A*, 408, 499
- Wolf, C., et al. 2004, *A&A*, 421, 913
- Wu, X.-B., Zhang, W., & Zhou, X. 2004, *Chinese, J. Astron. Astrophys.*, 4, 17
- York, D. G., et al. 2000, *AJ*, 120, 1579